

## **Remarks**

### Status of application

Claims 1-70 were examined and stand rejected in view of prior art. Applicant appreciates the courtesy of a telephone interview with the Examiner on October 18, 2007 to discuss Applicant's claims and the prior art references. The claims have been amended to further clarify Applicant's invention. Based on these amendments and the below remarks, reexamination and reconsideration are respectfully requested.

### Prior art rejections

Applicant's claims 48-70 stand rejected under 35 U.S.C. 102(e) as being anticipated by Graefe "Volcano - An Extensible and Parallel Query Evaluation System", IEEE Transactions on Knowledge and Data Engineering, Vol. 6 No 1, February 1994 (herein "Graefe '94") and claims 1-47 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Graefe '94 (above) in view of Graefe et al "Dynamic Query Evaluation Plans", Oregon Graduate Center, 1989, ACM (herein "Graefe '89").

Applicant appreciates the opportunity to discuss the pending claims and the prior art references with the Examiner. Based on these discussions, Applicant has amended the claims in a manner which Applicant believes clearly distinguishes Applicant's claimed invention from the prior art of record.

The focus of Applicant's invention is a specific improvement for optimizing query plans that may execute in parallel, by creating a schedule for each plan based on available resources and then comparing the alternative plans so as to choose the best plan.

Applicant's approach includes generating a plurality of different parallel plans and then adjusting these parallel plans to account for maximum configured threads available for executing the query. Furthermore, Applicant's claimed invention provides for generating a schedule for each parallel plan based on dependencies among operators of each parallel plan and available resources. The costs of each parallel plan are then determined based on the schedule and best plan selected for executing the query. Applicant's independent claims have been amended to bring these distinctive features to the forefront so as to more clearly distinguish Applicant's invention from the prior art. As will be shown below, Graefe '94 and Graefe '89, either alone or in combination fails to teach each and

every element set forth in Applicant's claims.

In Applicant's claimed invention, a plurality of parallel plans, structured as trees of operators, are generated for executing a given query. The operators included in such trees refer to constructs that reflect SQL operations, such as join operations, predicate evaluation mechanisms, scans (e.g., table or index), grouping operations, union operations, distinct operations and the like. In Applicant's system, these operations are examined to determine whether performing particular operations in parallel will yield improved performance in executing the query. However, the particular focus of Applicant's invention is finding out the best possible parallel plan given the resources which are available for executing the query.

Applicant's approach includes adjusting parallel plans (operator trees) based on applicable resource constraints, such as, for example, by the maximum number of configured threads available for executing the query. Although Graefe '94 describes an "exchange" iterator which is used for parallel execution and synchronization, this iterator only provides the mechanism for introducing parallelism into a query plan (Graefe '94 at page 130, Section VI, last paragraph). However, Graefe '94 does not teach or suggest that the degree of parallelism of a given operator tree is evaluated and adjusted if necessary based upon a configured maximum number of threads. Thus, it is not comparable to the specific teachings of Applicant's claimed invention of adjusting a parallel plan (operator tree) for executing a query based on maximum threads available for executing the query. Applicant's review of the balance of the Graefe '94 and Graefe '89 references finds no mention of an adjustment to ensure that an operator tree (plan) does not exceed the maximum number of configured threads as provided in Applicant's specification and claims.

In determining the best plan for executing a given query, Applicant's invention also provides for creation of a schedule for each parallel plan indicating the sequence for execution of operators of each parallel plan and associated costs in a manner that is not taught or suggested by the prior art. Parallel processing of queries provides a wide range of possibilities in how serial and parallel operations (including different levels of parallelism within a given operation) are combined. There is also another dimension of choices available in how a given operation is performed (e.g., table scan or index scan).

For example, a scan operation may be performed using a table scan or an index scan (as available). Accordingly, combining these dimensions of possibilities in all possible ways yields  $n$  different plans, where  $n$  may be a very large number.

For all of those possible plans, Applicant's system provides for computing a cost to determine which of the possible plans is the best plan. In determining the cost (elapsed time) for each possible plan, Applicant's system figures out how all the various operators are to be scheduled. Applicant's scheduling process includes examining dependencies among operators of the plans and resource constraints and costs. The schedule created for each parallel plan indicates when an operator would start and when would it stop, based on available resources (Applicant's specification, paragraph [136]; Figs. 7A-B at 705-712). Based on that information, that is the activation schedule of various operators, Applicant's system can determine the elapsed time for each particular plan (operator tree), and thus can determine which plan is in fact the best one. If the elapsed time of the current operator tree being considered is less than the minimum cost (i.e., elapsed time) of any previously evaluated operator tree, the current operator tree is saved as the best plan (best operator tree) and the elapsed time of the current plan is made the new minimum (Applicant's specification, paragraph [137]; Fig. 7B at 713-715). These features are not found in the prior art references.

Although Graefe '94 describes a "choose plan" operator, this choose plan operator simply calls a supporting function to decide which of several equivalent query plans should be utilized (Graefe '94, page 129, first column, second paragraph of Section V). The actual policy for selecting which plan to utilize is implemented by a supporting function, the details of which are not described. Thus, while the "choose plan" operator provides a mechanism to implement various policies for selecting from amongst alternative query plans, it does not itself implement any such policies. Therefore, it is not comparable to Applicant's claim limitations of generating a schedule for execution of each operator tree based upon dependencies among operators and available resources and evaluating costs of execution of alternative plans based on such schedule.

All told, the prior art references do not include teachings of adjusting operator trees (plans) based on threads and other resources available for executing a given query, nor do they describe anything comparable to Applicant's scheduling process which

selects generates a schedule for activation of operators of a particular plan based on dependencies among operators and available resources so as to determine the elapsed time for each particular plan (operator tree), and thus can determine which plan is the best plan for executing the query. As Graefe '94 and Graefe '89, either alone or in combination, do not include all the limitations of Applicant's claims it is respectfully submitted that Applicant's claims distinguish over the prior art and overcome any rejection under Section 102 and Section 103.

Any dependent claims not explicitly discussed are believed to be allowable by virtue of dependency from Applicant's independent claims, as discussed in detail above.

#### Conclusion

In view of the foregoing remarks and the amendment to the claims, it is believed that all claims are now in condition for allowance. Hence, it is respectfully requested that the application be passed to issue at an early date.

If for any reason the Examiner feels that a telephone conference would in any way expedite prosecution of the subject application, the Examiner is invited to telephone the undersigned at 925 465 0361.

Respectfully submitted,

Date: October 22, 2007

/G. Mack Riddle/

G. Mack Riddle; Reg. No. 55,572  
Attorney of Record

925 465 0361  
925 465 8143 FAX